

A Multiscale Retinex For Color Rendition and Dynamic Range Compression

Zia-ur Rahman[†]

College of William & Mary, Williamsburg, VA 23187

Daniel J. Jobson and Glenn A. Woodell

NASA Langley Research Center, Hampton, Virginia 23681

Abstract

The human vision system performs the tasks of dynamic range compression and color constancy almost effortlessly. The same tasks pose a very challenging problem for imaging systems whose dynamic range is restricted by either the dynamic response of film, in case of analog cameras, or by the analog-to-digital converters, in the case of digital cameras. The images thus formed are unable to encompass the wide dynamic range present in most natural scenes (often $> 500:1$). Whereas the human visual system is quite tolerant to spectral changes in lighting conditions, these strongly affect both the film response for analog cameras and the filter responses for digital cameras, leading to incorrect color formulation in the acquired image. Our multiscale retinex, based in part on Edwin Land's work on color constancy, provides a fast, simple, and automatic technique for simultaneous dynamic range compression and accurate color rendition. The retinex algorithm is non-linear, and global—output at a point is also a function of its surround—in extent. A comparison with conventional dynamic range compression techniques such as the application of point non-linearities, e.g. $\log(x, y)$, and global histogram equalization and/or modification shows that the multiscale retinex simultaneously provides the best dynamic range compression and color rendition. The applications of such an algorithm are many; from medical imaging to remote sensing; and from commercial photography to color transmission.

1. Introduction

Human perception excels at constructing a visual representation with vivid color and detail across wide ranging photometric levels caused by lighting variations. In addition human vision computes color so as to be relatively independent of spectral variations in illumination.¹ The images obtained with film and electronic cameras suffer, by comparison, from a loss in clarity of detail and color as light levels drop within shadows, or as distance from a lighting source increases. When the dynamic range of a scene exceeds the camera's dynamic range, there can be irrevocable loss of visual information at both extremes of the scene dynamic range. Improved fidelity of color images to human observation should, therefore, combine dynamic range compression, color constancy, and color and lightness rendition. In this paper we present our initial work in developing a technique, the multiscale retinex with color restoration (MSRCR), which achieves all these goals.

The idea of the retinex was conceived by Edwin Land^{2, 3, 4} as a model of the lightness and color perception of human vision. Subsequently Hurlbert^{5, 6}, and Hurlbert and Poggio⁷ studied the properties of the center/surround form of the retinex and other lightness theories and found a common mathematical foundation which possesses some excellent properties but cannot actually compute reflectance for arbitrary scenes. Certain scenes violate the "gray-world" assumption which requires that the average reflectances in the surround be equal in the three spectral color bands. For example, scenes that are dominated by one color—"monochromes"—clearly violate this assumption and are forced to gray (equal values in all spectral channels) by the retinex computation. Hurlbert⁵ further showed the lightness problem has a solution that has a center/surround spatial form. This suggests the possibility that the spatial opponency of the center/surround is a general solution to estimating relative reflectances for arbitrary lighting conditions. At the same time it is equally clear that human vision does not determine relative reflectance but rather a context dependent relative reflectance since surfaces in shadow do not appear to be the same lightness as the same surface when lit. Moore et al.^{8, 9} took up the retinex problem as a natural implementation for analog VLSI resistive networks and found that color rendition was dependent on scene content—some scenes worked well,

[†]Partially funded by NASA Langley Research Center contract #NAS1-19603 to Science and Technology Corporation.

others did not. These studies also pointed out the problems that occur with color Mach bands and the graying out of large uniform zones of color.

The MSRCR builds on the single scale retinex¹⁰ (SSR), and the multiscale retinex¹¹ (MSR). Both the SSR and the MSR provide very good dynamic range compression but suffer from the graying out which occurs in large areas of uniform color. Hence the overall color/lightness rendition can be poor depending upon the scene. The MSRCR alleviates this problem by using a color restoration function which controls the amount of color saturation for the final rendition. This function provides the color restoration that is needed with the dynamic range compression to approximate the performance of human vision with a computation that is quite automatic and reasonably simple. The MSRCR is extremely useful for enhancing 8-bit color images that suffer from lighting deficiencies commonly encountered in architectural interiors and exteriors, landscapes, and non-studio portraiture applications. Potential benefits for remote sensing applications are improved visibility of color and detail in shadows and low reflectance zones and the diminution of sun angle/atmospheric signal variations that can lead to more resilient and accurate multispectral classification.

2. Multiscale Center/Surround Retinex

The SSR^{10, 12, 13} is given by

$$R_i(x, y) = \log I_i(x, y) - \log [F(x, y) * I_i(x, y)] \quad (1)$$

where $R_i(x, y)$ is the retinex output, $I_i(x, y)$ is the image distribution in the i th color spectral band, "*" denotes the convolution operation, and $F(x, y)$ is the surround function,

$$F(x, y) = K e^{-(x^2+y^2)/c^2},$$

where c is the Gaussian surround space constant, or the scale, and K is selected such that


$$\iint F(x, y) dx dy = 1.$$

The MSR output is simply the weighted sum of the outputs of several SSRs with different scales. Mathematically,


$$R_{M_i}(x, y) = \sum_{n=1}^N w_n R_{n_i}(x, y), \quad (2)$$

where N is the number of scales, $R_{n_i}(x, y)$ is the i th component of the n th scale, $R_{M_i}(x, y)$ is the i th color component of the MSR output, and w_n is the weight associated with the n th scale. The number of scales is application dependent. However, after experimenting with one small scale and one large scale, the need for a third intermediate scale was immediately apparent in order to produce a graceful rendition without visible "halo" artifacts near strong edges. Experimentation shows that assigning equal weights to the scales is adequate for most applications, although a particular scale could be weighted more heavily if a particular feature needs to be enhanced. For instance, weighting the smallest scale heavily can be used to achieve the strongest dynamic range compression but leads to ungraceful edge artifacts and some graying of uniform color zones in the rendition.


To test whether the dynamic range compression of the MSR approaches that of human vision we use test SCENES not just test images, to facilitate the comparison between the processed image and direct observation. An example (Fig. 1) illustrates the complementary strengths and weaknesses of each scale taken separately and the strength of the multiscale synthesis. This image is representative of a number of test scenes (Fig. 2) where for conciseness we show only the multiscale result. The comparison of the unprocessed images to the perception of the scene produces some striking and unexpected results. Compared to recorded images, the color and detail are far more vivid for direct observation not only in shadowed regions, but also in the bright zones of the scene. This suggests that human




15 pixels




80 pixels



Original image



250 pixels



Multiscale

Figure 1: The components of the multiscale retinex which show their complementary information content. The smallest scale is strong on detail and dynamic range compression and weak on tonal and color rendition. The reverse is true for the largest spatial scale. The multiscale retinex combines the strengths of each scale and mitigates the weaknesses of each.




Figure 2: Examples of test SCENES processed with the multiscale retinex prior to color restoration. While color rendition of the left image is good, the other two are “grayed” to some extent. Dynamic range compression and tonal rendition are good for all and compare well with scene observation.

vision is perhaps doing more than just strong dynamic range compression and that enhancements beyond the MSR may be needed to capture the realism of direct viewing.

A sample of image data for surfaces in both sun and shadow indicates a dynamic range compression of 2 : 1 for the MSR compared to the 3 : 1 to 5 : 1 measured in our perceptual tests. For the SSR this value is 1.5 : 1 or less. These levels of dynamic range compression are for outdoor scenes which have shadows of large spatial extent. The much higher values of compression that occur for the human visual perception of mixed indoor/outdoor scenes are compared to retinex performance in Fig. 2 (right). The foreground orange book on the grayscale is compressed by approximately 5 : 1 for the MSR while compression for the SSR is only approximately 3 : 1 both relative to the bright building facade in the background. For this case, the compression of human vision is difficult to estimate since both the color and texture of the two surfaces are quite different. Our impression is that the MSR is approaching human vision’s performance but not quite reaching it.

The MSR performs well in terms of dynamic range compression but its performance on the pathological classes of images examined in previous SSR research¹⁰ (Fig. 3 middle row) still needs to be considered. These images represent a variety of regional and global gray-world violations and we can not expect the MSR to handle them effectively. We provide these results as a baseline for comparison with the color restoration which is developed next. All possess notable, and often serious, defects in color rendition. Since we want the MSR to be automatic, and the pathological images cannot be determined a priori, we developed an additional color computation which is universally applied to all post-retinex images to produce a general purpose computation.




Figure 3: Pathological “gray-world” violations are not handled well by the multiscale retinex alone (middle row) but are treated successfully when color restoration is added (lower row).

3. A Color Restoration Method for the Multi-scale Retinex

The general effect of retinex processing on images with regional or global gray-world violations is a “graying out” of the image either in specific regions or globally. This desaturation of color can, in some cases, be severe (Fig. 3 middle). Therefore we can consider the desired color computation as a color restoration, which should produce good color rendition for images with any degree of graying. More rarely, the gray-world violations can simply produce an unexpected color distortion (Fig. 3 top-left). Again we seek a simple computation which also handles these cases. In addition we would like for the correction to preserve a reasonable degree of color constancy since that is one of the basic motivations for the retinex. Color constancy is known to be imperfect in human visual perception, so some level of illuminant color dependency is acceptable provided it is much lower than the physical spectrophotometric variations. Ultimately this is a matter of image quality and color dependency is tolerable to the extent that the visual defect is not visually too strong.

Starting with the foundations of colorimetry¹⁴, the color space is transformed using

$$I'_i(x, y) = \frac{I_i(x, y)}{\sum_{j=1}^N I_j(x, y)}, i = 1, \dots, N. \quad (3)$$

The color restoration function $C(x, y)$ is then simply

$$C(x, y) \equiv C_i(x, y) = f\{I'_i(x, y)\},$$




Figure 4: Test SCENES illustrating dynamic range compression, color and tonal rendition, and automatic exposure correction. All processed images compare favorably with direct scene observation with the possible exception of leftmost image which is even lighter and clearer for observation. This scene has the widest dynamic range and suggests that even stronger dynamic range compression may be needed for this case.

where $f[\cdot]$ represents linearly or non-linearly normalized color space, and controls the saturation of the final rendition. The MSRCR is then given by

$$R_i(x, y) = C(x, y) \sum_{i=1}^N W_i (\log[I_i(x, y)] - \log[I_i(x, y) * F(x, y)]). \quad (4)$$

This form provides dynamic range compression, color and lightness constancy, and very good color rendition.

4. Selected Results for Diverse Test Cases

The test images presented here begin with some test scenes. We feel it is fundamental to refer the processed images back to the direct observation of scenes. This is necessary to establish that how well the computation represents a result that is "what you would have seen if you had been there". Clearly we cannot duplicate human vision's peripheral vision which spans almost 180° , but within the narrower angle of most image frames we would like to demonstrate that the computation achieves the clarity of color and detail in shadows, reasonable color constancy and lightness and color rendition that is present in direct observation of scenes. While we cannot yet test performance for scenes that go beyond 8-bit dynamic ranges, these results support the utility of the processing scheme for the enhancement of conventional 8-bit color images. The test scenes are given first (Figs. 4, 5) so that we can describe the degree to which the computation approaches human visual performance. All the test scene images after retinex processing are quite "true to life" compared with direct observation. We did not carefully match camera spatial resolution to observation so some difference in perceived detail is expected and observed. However overall color, lightness, and detail rendering for the multiscale retinex is a good approximation to human visual perception.

5. Discussion

The question which now arises is: What advantages does the MSRCR possess over traditional image enhancement techniques such as histogram equalization, non-linear transforms (gamma correction), and gain/offset manipulation?




Figure 5: Photographic examples further illustrating graceful dynamic range compression together with tonal and color rendition. The rightmost image shows the processing scheme handling saturated colors quite well and not distorting an image that is quite good in its original form.

Again the answer is based on experimental observation, rather than on theory. Each of the traditional techniques is well suited for a certain class of images, where the overall contrast is poor. They almost invariably fail where the image simultaneously contains very bright and very dark areas. They also fail to preserve the color when applied to images where the need for enhancement is not readily observable. The MSRCR successfully overcomes both these weaknesses of the traditional techniques. Figure 6 shows a comparison of the MSRCR with the traditional techniques for two natural scenes. The first contains a typical outdoor scene which has a sharp shadow across the frame. And the second is a *good* image which does not obviously need image enhancement. In both cases, the output of the MSRCR is either better than the original or as good. The same cannot be said of the traditional techniques.

The MSRCR can be applied *ex post facto* on 8-bit color images to provide image enhancement. The only problem arises when these images have been compressed using lossy methods. Not only does the MSRCR improve the dynamic range and color, it also enhances the compression artifacts which had been imperceptible before the application. Hence, the retinex is best applied prior to lossy image coding. One obvious advantage that the MSRCR provides for image compression is its ability to compress wide dynamic ranges to 8-bit, or less, per band color output, while preserving, and even enhancing, the details in the scene. The overall effect then is a significant reduction in the number of bits (especially in cases where the original color resolution is higher than 8-bit/band), required to transmit the original without a substantial loss in spatial resolution or contrast quality.

We have encountered many digital images in our testing which are either under- or overexposed. Apparently even with modern photographic auto-exposure controls, exposure errors can and do occur. An additional benefit of the MSRCR is its apparent capability for exposure correction. This is especially beneficial if it is performed before the image is recorded either on film or on disk.

6. Conclusions

The SSR provides a good mechanism for enhancing certain aspects of images and providing dynamic range compression. However, it is limited in its use because it can either provide good tonal rendition or dynamic range compression. The MSR comprised of three scales—small, intermediate, and large—overcomes this limitation

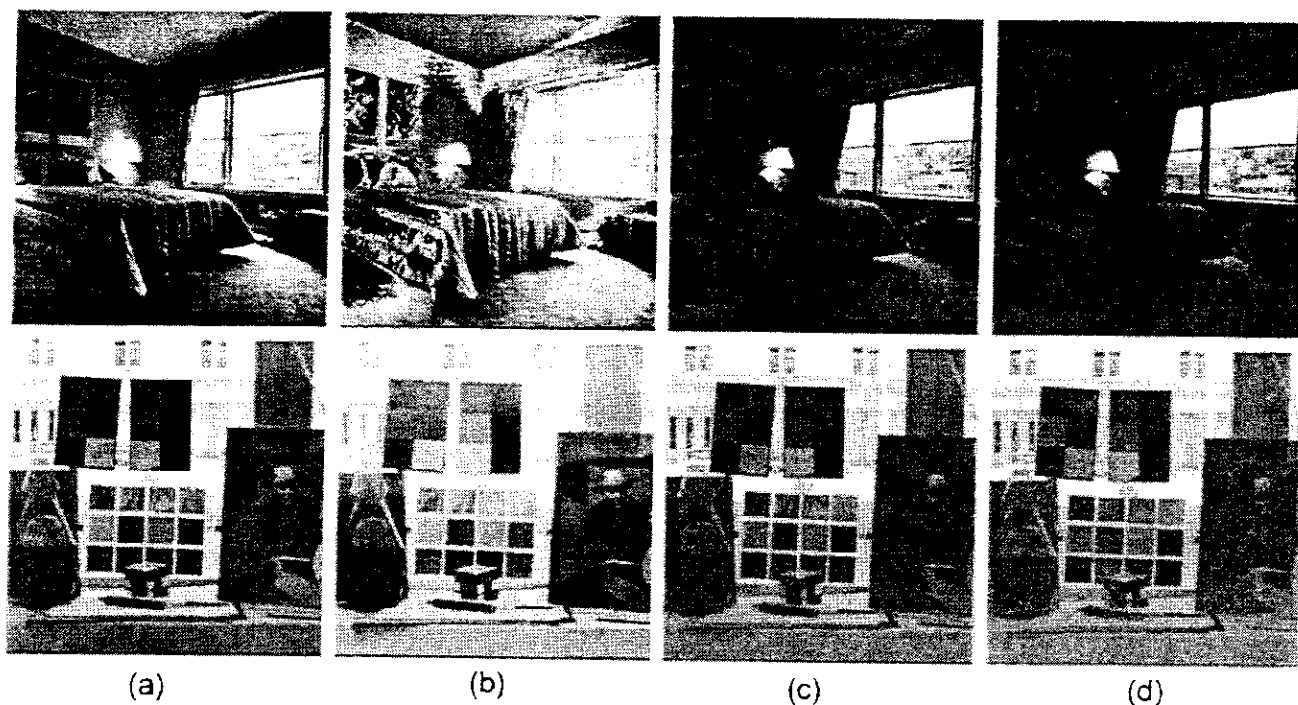


Figure 6: A comparison of image enhancement techniques: (a) MSRCR with 3 scales (b) Histogram Equalization, (c) Gamma correction, and (d) Gain/offset manipulation.

and was found to synthesize dynamic range compression, color constancy, and tonal rendition and produce results which compare favorably with human visual perception except for scenes which contain violations of the gray-world assumption. Even when the gray-world violations were not dramatic, some desaturation of color was found to occur. The MSRCR adds a color restoration scheme which produced good color rendition even for severe gray-world violations, but at the expense of a slight sacrifice in color constancy. While there is no firm theoretical or mathematical basis for proving the generality of the MSRCR, we have tested it successfully on numerous diverse scenes and images, including some known to contain severe gray-world violations.

REFERENCES

1. T. Cornsweet, *Visual Perception*. Orlando, FL: Academic Press, 1970.
2. E. Land, "Recent advances in retinex theory and some implications for cortical computations," *Proc. Nat. Acad. Sci.*, vol. 80, pp. 5163-5169, 1983.
3. E. Land, "An alternative technique for the computation of the designator in the retinex theory of color vision," *Proc. Nat. Acad. Sci.*, vol. 83, pp. 3078-3080, 1986.
4. E. Land, "Recent advances in retinex theory," *Vision Research*, vol. 26, no. 1, pp. 7-21, 1986.
5. A. C. Hurlbert, *The Computation of Color*. PhD thesis, Massachusetts Institute of Technology, September 1989.
6. A. C. Hurlbert, "Formal connections between lightness algorithms," *Journal of the Optical Society of America A*, vol. 3, pp. 1684-1693, 1986.
7. A. C. Hurlbert and T. Poggio, "Synthesizing a color algorithm from examples," *Science*, vol. 239, pp. 482-485, 1988.

8. A. Moore, J. Allman, and R. M. Goodman, "A real-time neural system for color constancy," *IEEE Transactions on Neural Networks*, vol. 2, pp. 237-247, March 1991.
9. A. Moore, G. Fox, J. Allman, and R. M. Goodman, "A VLSI neural network for color constancy," in *Advances in Neural Information Processing 3* (D. S. Touretzky and R. Lippman, eds.), pp. 370-376, San Mateo, CA: Morgan Kaufmann, 1991.
10. D. J. Jobson, Z. Rahman, and G. A. Woodell, "Properties and performance of a center/surround retinex," *IEEE Transactions on Image Processing*, 1996. Accepted for publication.
11. Z. Rahman, D. Jobson, and G. A. Woodell, "Multiscale retinex for color image enhancement," in *Proceedings of the IEEE International Conference on Image Processing*, IEEE, 1996.
12. Z. Rahman, "Properties of a center/surround Retinex Part One: Signal processing design," *NASA Contractor Report #198194*, 1995.
13. D. J. Jobson and G. A. Woodell, "Properties of a center/surround Retinex Part Two: Surround design," *NASA Technical Memorandum #110188*, 1995.
14. W. D. Wright, *The Measurement of Colour*. London: Hilger and Watts, second ed., 1958.